arXiv:1602.04399v1 [cs.DS] 13 Feb 2016

Linear Separability in Spatial Databases

Gilberto Gutiérrez* Pablo Pérez-Lantero† Claudio Torres*

February 16, 2016

Abstract

Given two point sets R and B in the plane, with cardinalities m and n, respectively, and each set stored in a separate R-tree, we present an algorithm to decide whether R and B are linearly separable. Our algorithm exploits the structure of the R-trees, loading into the main memory only relevant data, and runs in $O(m \log m + n \log n)$ time in the worst case. As experimental results, we implement the proposed algorithm and executed it on several real and synthetic point sets, showing that the percentage of nodes of the R-trees that are accessed and the memory usage are low in these cases. We also present an algorithm to compute the convex hull of n planar points given in an R-tree, running in $O(n \log n)$ time in the worst case.

1 Introduction

Let R be a finite set of red points and B a finite set of blue points in the plane, with cardinalities m and n, respectively. We say that R and B have *linear separability*, or that are *linearly separable*, if there exists a line such that: the elements of R belong to one of the halfplanes bounded by the line, the elements of B belong to the other halfplane, and if the line contains points from $R \cup B$, then it contains points from exactly one between R and B. It is well known that R and B are linearly separable if and only if the convex hulls $conv(R)$ and $conv(B)$ have an empty intersection, where $conv(X)$ denotes the convex hull of $X \subset \mathbb{R}^2$. Because of this, deciding whether R and B are linearly separable, and in the positive case finding such a separating line, can be done within the following steps: compute $conv(R)$ and $conv(B)$ in times $O(m \log m)$ and $O(n \log n)$, respectively; and test whether the intersection of $conv(R)$ and $conv(B)$ is empty in time $O(m+n)$ [18]. In fact, if $conv(R)$ and $conv(B)$ has an empty intersection, then there exists a separating line containing an edge of $conv(R)$ or $conv(B)$ [5]. Another way of deciding linear separability of R and B, without computing the convex hulls, is to formulate the separation problem as a linear program (LP) in two variables and $m+n$ constraints, and use the algorithm of Megiddo [14] or the algorithm of Dyer [6] to solve the LP in linear $O(m+n)$ time in the worst case. Other more practical option, due to the big hidden constants in the $O(m+n)$ running times of these two algorithms, is to use the simpler randomized algorithm of Seidel [16] that solves the LP in expected $O(m+n)$ time.

*Computer Science and Information Technologies Department, Universidad del Bío-Bío, Chillán, Chile. ggutierr@ubiobio.cl; ctorresf@alumnos.ubiobio.cl.

†Departamento de Matemática y Ciencia de la Computación, Universidad de Santiago de Chile (USACH), Santiago, Chile. pablo.perez@usach.cl.

In this paper, we study the problem of deciding whether R and B are linearly separable, and at the same time returning a separating line when the answer is positive, in the context of the spatial databases: We are given as input an R-tree with the points of R and a second R-tree with the points of B. The R-tree is a secondary-storage, tree-like height-balanced data structure designed for the dynamic indexation of a set of dimensional geometric objects [9, 13], and it is considered an standard in the context of the spatial databases. See Section 1.1 for further details. Because of the high data volume of the R-trees, loading from the R-trees all points of $R \cup B$ and running a known algorithm for testing linear separability of R and B is considered impractical in this context. Then, we aim to design an efficient algorithm working directly with the R-tree structure and able of loading from the R-trees only the relevant data. Most of the algorithms in this context, apart of producing the correct answer, aim to minimize both the running time and the number of nodes read from the R-trees, since each node is implemented as a disk page.

The spatial databases (SDB) represent an important aid for geographical information systems (GIS) to manage large amounts of data. However, SDB require the design of new data structures, spatial access methods, query languages, and algorithms to manage spatial information. In this sense, several algorithm have been designed for spatial queries such as the window query, the intersection query, the nearest neighbor, and the spatial join [7, 17]. Many of these queries are problems that were first tackled in the field of the computational geometry, where it is assumed that all spatial objects fit into the main memory, and later, these problems were faced in the field of the SDB. Following this path, several algorithms have been proposed considering that objects are stored in a multidimensional structure, in most cases an R-tree [9]. For example, Corral et al. [4] and Hjaltason et al. [12] presented several algorithms that solve the k-pairs ($k \geq 1$) of nearest neighbors between two sets, Roussopoulos et al. [15] showed an algorithm to find the nearest neighbor to a given point, Gutiérrez et al. [8] showed how to find a largest rectangle containing a query object and no point stored at an input R-tree, and Böhm and Kriegel [2] described methods for computing the convex hull of point sets stored in hierarchical index structures such as the R-tree or its variants. Among the geometric problems in spatial databases, this later work is well related to the results of this paper since linear the separability of two point sets can be decided by computing the convex hull of each set, and querying the disjointness of the convex hulls. It is worth noting that Böhm and Kriegel's algorithms do not exploit particular properties of R-trees such as the fact that the node regions are minimum bounding rectangles. We explicitly exploit such a property in the algorithm that we propose for deciding the linear separability of two point sets in the plane. Furthermore, in many cases our algorithm can decide the linear separability without computing such convex hulls.

The linear separability of two point sets, in either the plane or higher dimensions where a hyperplane separates the two point sets, is a concept well used in data mining and classification problems [1, 11]. In this setting, where each point set represents the data of one class, the support vector machines (SVM) are a very robust methodology for inference [1]. The SVM is the hyperplane that separates the point sets and maximizes the minimum distance from the points to it.

The results of this paper are the following ones:

(1) We present an algorithm working directly with the R-trees of R and B, able of deciding whether R and B are linearly separable, and in the positive case returning a separating line. In each step, it loads in the main memory data of only one level for each of the R-trees, and before descending in the R-trees to the next levels, to finally end at the leaf nodes, it verifies whether the gathered information is enough to decide the separability condition of R and B.

The asymptotic running time in the worst case is $O(m \log m + n \log n)$.

(2) The techniques used in the separability testing algorithm can be extended to design an algorithm that computes the convex hull of a finite planar point set given as input in an R-tree. If n denotes the number of input points, the asymptotic running time is $O(n \log n)$.

(3) We implement the separability testing algorithm and executed it on several real and synthetic colored point sets, showing that in both cases the number of nodes of the R-trees that are accessed by the algorithm, and the amount of memory used, are low for these point sets. To generate synthetic point sets, we consider parameters such as the number of points to generate, the distribution of the points (uniform or Gaussian), and other parameters to define the positions of the minimum bounding rectangles of the red and blue points, respectively.

Notation: Given a set $X \subset \mathbb{R}^2$, let $\mathsf{MBR}(X)$ denote the Minimum Bounding Rectangle (MBR) of X, which is the minimum-area rectangle that contains X. Every rectangle is considered axis-aligned in this paper. Note that each of the four sides of $\mathsf{MBR}(X)$ contains at least one point of X. Extending the notation, if Y is a set of subsets of $\mathbb{R}^2$ (e.g. a set of rectangles), then $conv(Y)$ (resp. $\mathsf{MBR}(Y)$) denotes the convex hull (resp. MBR) of all points contained in some element of Y, and $conv(Y \cup X)$ denotes the convex hull of the union of the points of $conv(Y)$ and $conv(X)$.

Outline: We continue this section by describing the R-tree data structure in Section 1.1, and the idea of our separability algorithm in Section 1.2. In Section 2, we present the concepts and geometric properties that our algorithm uses. Later, in Section 4, we present our linear separability test algorithm, together with the algorithm to compute the convex hull of a point set given in an R-tree. In Section 5, we show the experimentation results. Finally, in Section 6, we state the conclusions and further research.

1.1 The R-tree

An R-tree is a generalization of the B^+-trees designed for the dynamic indexation of a set of k-dimensional geometric objects [13]. It is a hierarchical, height balanced multidimensional data structure, designed to be used in secondary storage. In inner levels the indexed objects are represented by the k-dimensional Minimum Bounding Rectangles (MBRs), which bound their children. In this paper, we focus on two dimensions, therefore each MBR is an axis-aligned rectangle, represented only by its bottom-left and top-right vertices. By using the MBRs instead of the exact geometrical representations of the objects, its representational complexity is reduced to two points where the most important features of the spatial object (position and extent) are maintained. The MBR is an approximation widely employed, and the R-trees belong to the category of data-driven access methods, since their structure adapts itself to the MBRs distribution in the space.

An R-tree for a finite point set $X \subset \mathbb{R}^2$ satisfies the following properties. The leaves are on the same level, and each leaf node contains indexed points of X. Every internal node contains entries of the form `(MBR,ref)`, where `ref` is a pointer to the child of the entry, and `MBR` is the minimum bounding rectangle of the `MBR`'s (or the points if the child is a leaf node) contained in the entries of this child. An R-tree has the property that every node, except possibly the root, contains between m and M entries, where $2 \leq \mathsf{m} \leq \lceil \mathsf{M}/2 \rceil$. The root node contains at least two children nodes. We

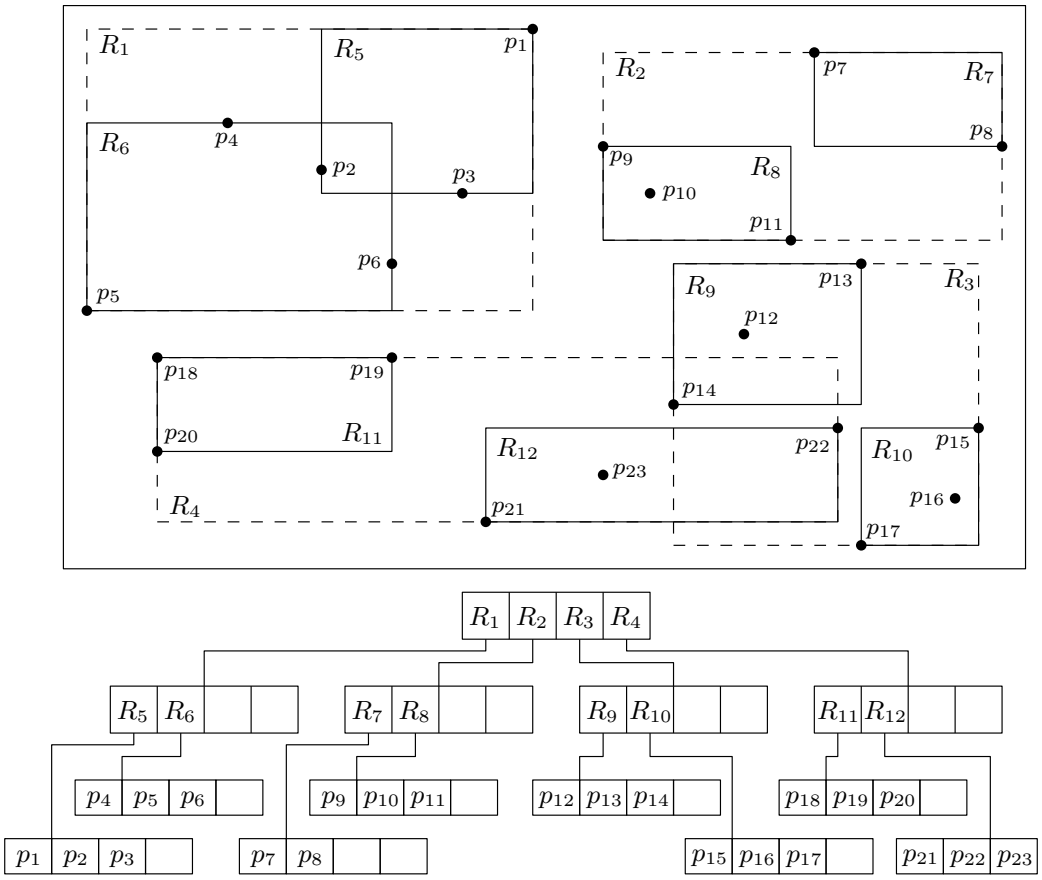

Figure 1: An R-tree for the point set $\{p_1, p_2, \ldots, p_{23}\}$ (picture based on one from [8]).

will refer to the MBR of an entry just as the rectangle of the entry, and to the MBR's of the entries of a node just as the rectangles of the node. For every entry (i.e. rectangle) of a node, the rectangles of the child node of the entry (i.e. the child rectangles) are not necessarily pairwise disjoint, so they can overlap between them. Furthermore, points of X could be covered by different rectangles of the R-tree, although each point of X appears only once in the leaf nodes. All nodes of an R-tree are implemented as disk pages. We consider that the leaf nodes are at level h and the root is at level 0, where h is the height of the R-tree.

Figure 1 depicts an R-tree. Dotted lines denote the rectangles of the entries at the root node. The rectangles with solid lines are the rectangles in the entries of nodes parent of the leaves. Finally, the points are the indexed objects in the leaves of the R-tree.

1.2 Idea of our algorithm

The general idea of our algorithm is based in the following observation: Suppose that have loaded in the main memory a set N_R of rectangles of the R-tree of R so that they all cover R, and a set of rectangles N_B of the R-tree of B so that they all cover B. If the convex hull $conv(N_R)$ of the rectangles of N_R does not intersect the convex hull $conv(N_B)$ of the rectangles of N_B, then $conv(R)$ and $conv(B)$ are disjoint since $conv(R) \subseteq conv(N_R)$ and $conv(B) \subseteq conv(N_B)$, and R and B are

hence linearly separable.

A more concrete idea is the following: According to the relative positions of $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$, we choose from $\mathsf{MBR}(R)$ a set V_R of at most three vertices, and a similar set V_B from $\mathsf{MBR}(B)$. The idea of choosing V_R and V_B is that $conv(R)$ and $conv(B)$ are disjoint if and only if $conv(R \cup V_R)$ and $conv(B \cup V_B)$ are disjoint. We start with N_R being the set of the rectangles stored in the root node of the R-tree of R, and N_B being the set of the rectangles stored at the root node of the R-tree of B. Then, we iterate as follows: If $conv(N_R \cup V_R)$ and $conv(N_B \cup V_B)$ are disjoint, then we report a 'yes' answer and build a separating line. Otherwise, for each rectangle of N_R we take the region of points that can be ensured to belong to $conv(R \cup V_R)$ and form the set of regions N'_R. A similar set N'_B is formed from N_B. If $conv(N'_R \cup V_R)$ and $conv(N'_B \cup V_B)$ are not disjoint, then we report a 'no' answer. Otherwise, we filter the set N_R so that the new N_R contains only the rectangles (or points) that are relevant to decide the linear separation of R and B, and 'refine' $conv(N_R \cup V_R)$ by replacing each rectangle in N_R by its respective child rectangles (or points) in the R-tree. We do a similar procedure with N_B and continue the iteration. If at some point in the iteration both N_R and N_B consist of only points, then the answer is given by the intersection condition of $conv(N_R \cup V_R)$ and $conv(N_B \cup V_B)$.

It is worth noting that we test the linear separability condition via computing the convex hulls (or approximations of the convex hulls, e.g. $conv(N_R \cup V_R)$) of the two point sets. We do not use any asymptotic-faster linear-time LP separability testing algorithm [6, 14, 16] since in this case the process of filtering rectangles is more expensive in time: We can discard a rectangle if we can ensure that it is contained in the convex hull, and it can be done in logarithmic time, as we will see later in the paper. Otherwise, if we do not compute the convex hull, as it happens if we use any of the linear-time LP separability testing algorithms, then to decide whether a rectangle can be discarded we should consider the relative position of the rectangle with respect to the other rectangles, and this is much more expensive than determining whether the rectangle is inside a convex hull.

2 Preliminaries

Considering $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$, we make the following definitions:

- We say that $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a *corner* intersection if each rectangle contains exactly one vertex of the other one (see Figure 2a), or one of the rectangles is contained in the other and they share exactly one vertex.
- We say that $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a *side* intersection if one of the rectangles contains exactly two vertices of the other one (see Figure 2b), and is not contained within it.
- We say that $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a *containment* intersection if the interior of one of the rectangles contains from the other rectangle the four vertices (see Figure 2c), or two adjacent vertices with the other two ones contained in the boundary.
- We say that $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a *piercing* intersection if their interiors are not disjoint, and no rectangle contains in the interior a vertex of the other one (see Figure 2d).

Up to symmetry, we assume without loss of generality throughout this paper that the relative positions of $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$, when they intersect, are those shown in Figure 2. Observe that

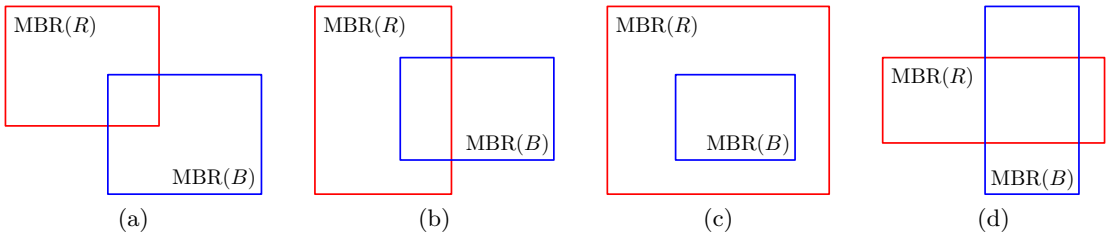

Figure 2: $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a: (a) corner intersection. (b) side intersection. (c) containment intersection. (d) piercing intersection.

if $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a piercing intersection, then R and B are not linearly separable. This is because their convex hulls have a non-empty intersection since each side of $\mathsf{MBR}(R)$ contains a red point and each side of $\mathsf{MBR}(B)$ contains a blue point. Furthermore, the piercing intersection definition includes the case where $\mathsf{MBR}(R) = \mathsf{MBR}(B)$. In the other cases of intersection (corner, side, and containment) the linear separation condition cannot be directly deduced (see for examples Figure 3). In the trivial case where $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ do not intersect, R and B are linearly separable with a vertical or horizontal line.

Our algorithm to decide whether R and B are linearly separable starts by detecting the type of intersection between $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$. If they do not intersect, then we report a 'yes'. If they have a piercing intersection, then we report a 'no'. Otherwise, if the intersection is of type corner, side, or containment, then we need to elaborate a procedure that gives the correct answer and at the same time returns a line separating R and B if the answer is 'yes'. Suppose that $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a containment intersection, with $\mathsf{MBR}(B)$ contained in $\mathsf{MBR}(R)$. In this case, the rectangles $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B \cup \{v\})$ have a corner intersection, where v is any of the four vertices of $\mathsf{MBR}(R)$. Furthermore, if $conv(R)$ and $conv(B)$ are disjoint, then $conv(R)$ and $conv(B \cup \{v\})$ are also disjoint for some vertex v of $\mathsf{MBR}(R)$. Conversely, if $conv(R)$ and $conv(B)$ are not disjoint, then $conv(R)$ and $conv(B \cup \{v\})$ will not be disjoint for every vertex v of $\mathsf{MBR}(R)$ since $conv(B)$ is contained in $conv(B \cup \{v\})$. Hence, when $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a containment intersection, we can test whether R and B are linearly separable by calling four times the test for the case where $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a corner intersection. That is, the answer is 'yes' if and only if R and $B \cup \{v\}$ are linearly separable for at least one vertex v of R. Because of this reduction, we will consider in the following only intersections of type corner or side.

Definition 1. *If* $\mathsf{MBR}(R)$ *and* $\mathsf{MBR}(B)$ *have a corner intersection, then let* V_R *denote the set of the top-right, top-left, and bottom-left vertices of* $\mathsf{MBR}(R)$*, and* V_B *denote the set of the top-right, bottom-left, and bottom-right vertices of* $\mathsf{MBR}(B)$*. If* $\mathsf{MBR}(R)$ *and* $\mathsf{MBR}(B)$ *have a side intersection, then let* V_R *denote the set of the top-left and bottom-left vertices of* $\mathsf{MBR}(R)$*, and* V_B *denote the set of the top-right and bottom-right vertices of* $\mathsf{MBR}(B)$ *(see figures 4a and 4b).*

The idea behind the definition of V_R and V_B is given in the next lemmas:

Lemma 2. *The convex hulls* $conv(R)$ *and* $conv(B)$ *are disjoint if and only if the convex hulls* $conv(R \cup V_R)$ *and* $conv(B \cup V_B)$ *are disjoint.*

Proof. If $conv(R \cup V_R)$ and $conv(B \cup V_B)$ are disjoint, then trivially $conv(R)$ and $conv(B)$ are also disjoint because $conv(R) \subseteq conv(R \cup V_R)$ and $conv(B) \subseteq conv(B \cup V_B)$. Then, suppose that $conv(R)$

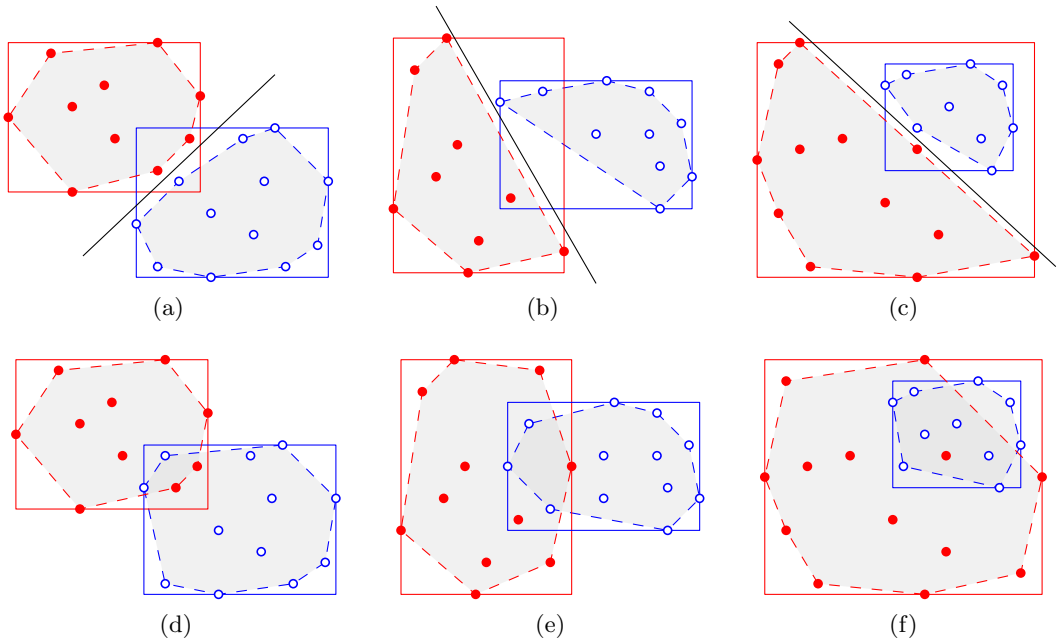

Figure 3: (a,b,c) Corner, side, and containment intersections of $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$, where R and B have linear separability. (d,e,f) Corner, side, and containment intersections of $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$, where R and B are not linearly separable. In each picture, the red points are represented as solid dots, and the blue points as tiny disks.

and $conv(B)$ are disjoint (see Figure 4c). Let p be a point of $conv(R \cup V_R) \setminus conv(R)$. If p does not belong to $\mathsf{MBR}(B)$, then p is not contained in $conv(B \cup V_B)$. Then, assume that p belongs to the intersection $\mathsf{MBR}(R) \cap \mathsf{MBR}(B)$. Let h_p be the halfline with apex at p, and oriented rightwards or downwards, such that h_p contains a point q of $conv(R) \cap \mathsf{MBR}(B)$. Note that h_p always exists given p, and assume without loss of generality that h_p is horizontal. The case where h_p is vertical (which appears only in the case of a corner intersection) is analogous. Since q does not belong to $conv(B)$ because $conv(R)$ and $conv(B)$ are disjoint, and q is in the interior of $\mathsf{MBR}(R) \cap \mathsf{MBR}(B)$, q is both to the left and above of $conv(B)$ in the case of a corner intersection, and to the left of $conv(B)$ in the case of a side intersection. Then, q does not belong to $conv(B \cup V_B) \setminus conv(B)$ because of the definition of V_B. Hence, p is not in $conv(B \cup V_B) \setminus conv(B)$ because of the convexity of $conv(B)$ and that p is to left of q. Similar symmetric arguments show that if a point p' belongs to $conv(B \cup V_B) \setminus conv(B)$, then p' is not in $conv(R \cup V_R) \setminus conv(R)$. All of these observations imply that $conv(R \cup V_R)$ and $conv(B \cup V_B)$ are disjoint. □

Lemma 3. *The convex hulls $conv(R \cup V_R)$ and $conv(B \cup V_B)$ have a non-empty intersection if and only if one of them contains a vertex of the other one. Furthermore, the vertices of $conv(R \cup V_R)$ contained in $conv(B \cup V_B)$ are all consecutive, and the vertices of $conv(B \cup V_B)$ contained in $conv(R \cup V_R)$ are all consecutive.*

Proof. If one of $conv(R \cup V_R)$ and $conv(B \cup V_B)$ contains a vertex of the other one, then they are not disjoint. Suppose now that $conv(R \cup V_R)$ and $conv(B \cup V_B)$ are not disjoint. Let X be a set of red points and Y a set of blue points. If $conv(X)$ and $conv(Y)$ are not disjoint and neither of them

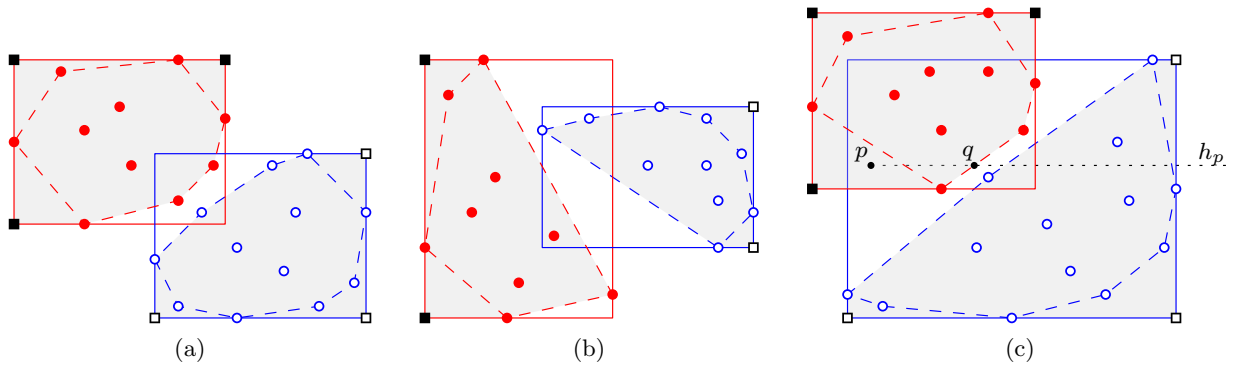

Figure 4: The sets of vertices V_R and V_B for: (a) corner intersection; (b) side intersection. The vertices of V_R are denoted as filled squares, and the vertices of V_B as empty squares. (c) Proof of Lemma 2. In each picture, the convex hulls $conv(R \cup V_R)$ and $conv(B \cup V_B)$ are denoted as shaded regions.

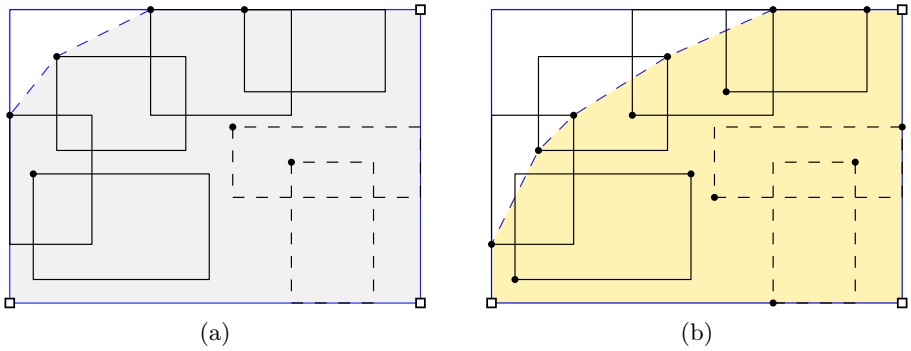

Figure 5: A set N_B of rectangles from the R-tree of B, in the case where $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a corner intersection. (a) The optimistic convex hull $\mathsf{opt}(N_B)$. (b) The pessimistic convex hull $\mathsf{pess}(N_B)$. The rectangles in dashed lines can be removed from N_B since they are contained in $\mathsf{pess}(N_B)$.

contains a vertex of the other one (see e.g. the piercing intersection of Figure 2d), then the convex hull $conv(X \cup Y)$ contains at least four bichromatic edges (i.e. edges connecting points of different colors). For $X = R \cup V_R$ and $Y = B \cup V_B$, $conv(X \cup Y)$ contains only two bichromatic edges, given the relative positions of $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ and the definitions of V_R and V_B. Hence, one of $conv(R \cup V_R)$ and $conv(B \cup V_B)$ must contain a vertex of the other one. Furthermore, the fact that $conv(X \cup Y)$ contains only two bichromatic edges implies the second part of the lemma. □

3 Optimistic and pessimistic Convex Hulls

The idea in this section is the following: Suppose that we have loaded a set of rectangles from the R-tree of R, and a set of rectangles from the R-tree of B. We explain a way of determining from these two sets of rectangles whether we have enough information to decide that $conv(R \cup V_R)$ and $conv(B \cup V_B)$ are disjoint or that they are not disjoint, without going deeper in the R-tree loading more rectangles or points.

Definition 4. *Let N_R be a set of rectangles from the R-tree of R, and N_B a set of rectangles from the R-tree of B, such that the following properties are satisfied:*

(1) $conv(R \cup V_R) \subseteq conv(N_R \cup R)$.

(2) $conv(B \cup V_B) \subseteq conv(N_B \cup B)$.

For examples of Definition 4, refer to Figure 5 and Figure 6. Since a point can be seen as a rectangle of null perimeter, we extend the definitions of N_R and N_B so that these sets can be made of points.

Definition 5. *Given the sets N_R and N_B, the optimistic convex hull of $R \cup V_R$ is the set* $\mathsf{opt}(N_R) = conv(N_R \cup V_R)$ *which contains* $conv(R \cup V_R)$*, and the optimistic convex hull of $B \cup V_B$ is set* $\mathsf{opt}(N_B) = conv(N_B \cup V_B)$ *which contains* $conv(B \cup V_B)$ *(see Figure 5a and Figure 6a).*

The idea of defining the optimistic convex hulls is the following observation: If $\mathsf{opt}(N_R)$ and $\mathsf{opt}(N_B)$ are disjoint, then we can ensure that $conv(R \cup V_R)$ and $conv(B \cup V_B)$ are disjoint, and then R and B are linearly separable because of Lemma 2. Furthermore, $\mathsf{opt}(N_R)$ and $\mathsf{opt}(N_B)$ are approximations to $conv(R \cup V_R)$ and $conv(B \cup V_B)$, and to compute them we do not need the points covered by their rectangles, which are located in the leaves of the R-trees.

We also need a method for determining from N_R and N_B whether there is enough information to decide that $conv(R \cup V_R)$ and $conv(B \cup V_B)$ are not disjoint. This is explained in what follows.

Let N be a rectangle: $\mathtt{NE}(N)$ denotes the north-east triangle of N, that is, the subset of points of N in or above the diagonal connecting the top-left and bottom-right vertices. Similarly, $\mathtt{NW}(N)$ denotes the subset of points of N in or above the diagonal connecting the top-right and bottom-left vertices; $\mathtt{SE}(N)$ denotes the subset of points of N in or below the diagonal connecting the top-right and bottom-left vertices; and $\mathtt{SW}(N)$ denotes the subset of points of N in or below the diagonal connecting the top-left and bottom-right vertices.

Suppose that $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a corner intersection, and let N be a rectangle of N_R. Since in our model of R-tree every rectangle is a minimum bounding rectangle and thus contains points of the represented point set in every side, the set $\mathtt{NW}(N)$ is contained in the convex hull $conv(R \cup V_R)$. Similarly, if N is a rectangle of N_B, then the set $\mathtt{SE}(N)$ is contained in the convex hull $conv(B \cup V_B)$. We use these observations to define the pessimistic convex hulls, which are always contained in our goal $conv(R \cup V_R)$ and $conv(B \cup V_B)$.

Definition 6. *Let R and B be red and blue point sets such that* $\mathsf{MBR}(R)$ *and* $\mathsf{MBR}(B)$ *have a corner intersection. The pessimistic convex hull of $R \cup V_R$ is the set* $\mathsf{pess}(N_R) = conv(N'_R \cup V_R)$ *which is contained in* $conv(R \cup V_R)$*, where* $N'_R = \{\mathtt{NW}(N) \mid N \in N_R\}$*. The pessimistic convex hull of $B \cup V_B$ is the set* $\mathsf{pess}(N_B) = conv(N'_B \cup V_B)$ *which is contained in* $conv(B \cup V_B)$*, where* $N'_B = \{\mathtt{SE}(N) \mid N \in N_B\}$ *(see Figure 5b).*

Definition 7. *Let R and B be red and blue point sets such that* $\mathsf{MBR}(R)$ *and* $\mathsf{MBR}(B)$ *have a side intersection. The pessimistic convex hull of $R \cup V_R$ is the set* $\mathsf{pess}(N_R) = conv(N'_R \cup V_R \cup \{u_1\}) \cap conv(N''_R \cup V_R \cup \{u_2\})$*, where* $N'_R = \{\mathtt{NW}(N) \mid N \in N_R\}$*,* $N''_R = \{\mathtt{SW}(N) \mid N \in N_R\}$*, and u_1 and u_2 are the top-right and bottom-right vertices of* $\mathsf{MBR}(R)$*, respectively. The pessimistic convex hull of $B \cup V_B$ is the set* $\mathsf{pess}(N_B) = conv(N'_B \cup V_B \cup \{v_1\}) \cap conv(N''_B \cup V_B \cup \{v_2\})$*, where* $N'_B = \{\mathtt{SE}(N) \mid N \in N_B\}$*,* $N''_B = \{\mathtt{NE}(N) \mid N \in N_B\}$*, and v_1 and v_2 are the bottom-left and top-left vertices of* $\mathsf{MBR}(B)$*, respectively. (see Figure 6).*

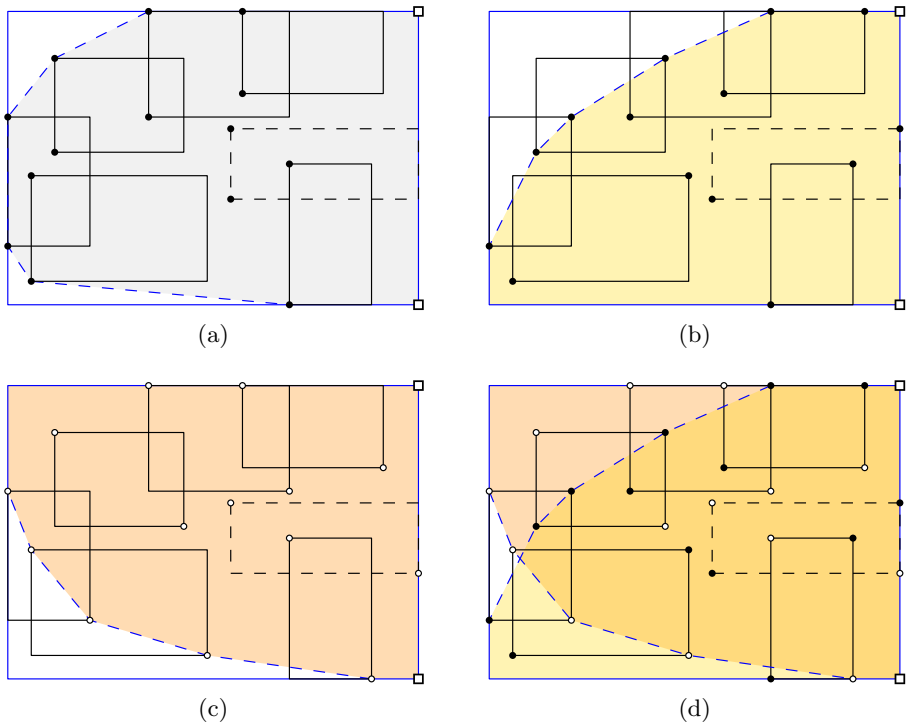

Figure 6: A set N_B of rectangles from the R-tree of B, when $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a side intersection. (a) The optimistic convex hull $\mathsf{opt}(N_B)$. (b) The convex hull $conv(N'_B \cup V_B \cup \{v_1\})$. (c) The convex hull $conv(N''_B \cup V_B \cup \{v_2\})$. (d) The pessimistic convex hull $\mathsf{pess}(N_B) = conv(N'_B \cup V_B \cup \{v_1\}) \cap conv(N''_B \cup V_B \cup \{v_2\})$. The rectangle in dashed lines can be removed from N_B since it is contained in $\mathsf{pess}(N_B)$.

When $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a corner intersection, the facts $\mathsf{pess}(N_R) \subseteq conv(R \cup V_R)$ and $\mathsf{pess}(N_B) \subseteq conv(B \cup V_B)$ are clear. In such a case of intersection, if $\mathsf{pess}(N_R)$ and $\mathsf{pess}(N_B)$ are not disjoint, then we can ensure that $conv(R \cup V_R)$ and $conv(B \cup V_B)$ are not disjoint, and then R and B are not linearly separable because of Lemma 2. When $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a side intersection, the same facts are proved in the following lemma, and then we can ensure that $conv(R \cup V_R)$ and $conv(B \cup V_B)$ are not disjoint if $\mathsf{pess}(N_R)$ and $\mathsf{pess}(N_B)$ are not.

Lemma 8. *Let R and B be red and blue point sets such that $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a side intersection. The pessimistic convex hulls $\mathsf{pess}(N_R)$ and $\mathsf{pess}(N_B)$ are contained in $conv(R \cup V_R)$ and $conv(B \cup V_B)$, respectively.*

Proof. We will prove that $\mathsf{pess}(N_B) \subseteq conv(B \cup V_B)$. The arguments to prove $\mathsf{pess}(N_R) \subseteq conv(R \cup V_R)$ are analogous. From the rectangles of N_B whose top sides are aligned with the top side of $\mathsf{MBR}(B)$, let t_1 a left-most of the top-right vertices (see Figure 7). From the rectangles of N_B whose bottom sides are aligned with the bottom side of $\mathsf{MBR}(B)$, let b_1 a left-most of the bottom-right vertices. From the rectangles of N_B whose left sides are aligned with the left side of $\mathsf{MBR}(B)$, let ℓ_1 be a bottom-most of the top-left vertices, and ℓ_2 a top-most of the bottom-left vertices. Observe that ℓ_2 and t_1 are vertices of $conv(N'_B \cup V_B \cup \{v_1\})$, and that b_1 and ℓ_1 are vertices of $conv(N''_B \cup V_B \cup \{v_2\})$. Let U denote the path connecting ℓ_2 with t_1 along the boundary of

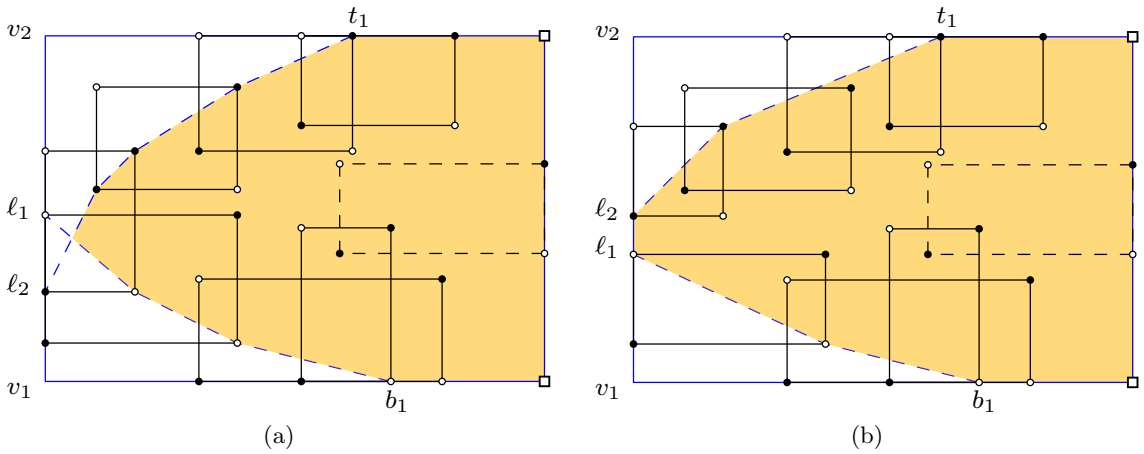

Figure 7: Proof of Lemma 8.

$conv(N'_B \cup V_B \cup \{v_1\})$ and clockwise, and L denote the path connecting b_1 with ℓ_1 along the boundary of $conv(N''_B \cup V_B \cup \{v_2\})$ and clockwise. Note that there are the following blue points of B: a left-most blue point t in the segment connecting v_2 and t_1, a left-most blue point b in the segment connecting v_1 and b_1, and both a top-most blue point ℓ and a bottom-most blue point ℓ' (possibly equal to ℓ) in the edge connecting v_1 and v_2. By the definitions of ℓ_1 and ℓ_2, we have that ℓ belongs to the segment connecting ℓ_2 and v_2, and ℓ' belongs to the segment connecting ℓ_1 and v_1. The key observation is that b, ℓ', ℓ, and t are all vertices of $conv(B \cup V_B)$. Furthermore, the clockwise path along the boundary of $conv(B \cup V_B)$ that connects b and ℓ'' is below the path L, and the similar path that connects ℓ and t is above the path U. All of these observations imply that $\mathsf{pess}(N_B) \subseteq conv(B \cup V_B)$. □

4 Algorithms

In this section, we present our separability testing algorithm for R and B in the R-tree model. We start by presenting the ingredient algorithms for the separability algorithm: computation of the optimistic and pessimistic convex hulls (Section 4.1), deciding whether the convex hulls are disjoint and finding a separating line in the positive case (Section 4.2), and filtering the sets of rectangles N_R and N_B (Section 4.3). Then, we show the separability algorithm (Section 4.4). Finally, we show how the techniques to previous computations can be used to compute the convex hull of a point set given in an R-tree (Section 4.5).

4.1 Convex hulls computation

We explain how to compute the optimistic and pessimistic convex hulls for N_B in both cases of intersections of $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$: corner and side. The algorithms to compute these convex hulls for N_R are analogous by symmetry.

Let t_1 (resp. t_2) be a left-most vertex from the top-right (resp. top-left) vertices of the rectangles of N_B whose top sides are aligned with the top side of $\mathsf{MBR}(B)$; b_1 (resp. b_2) a left-most vertex

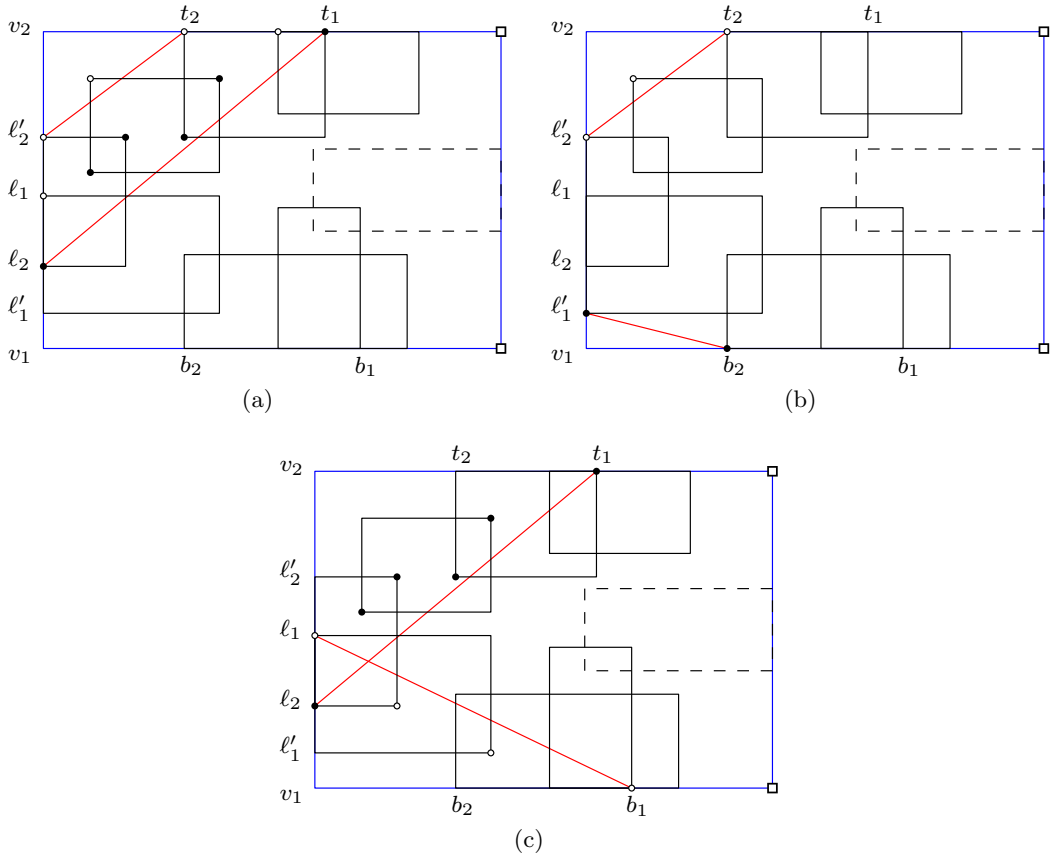

Figure 8: Algorithm to compute $\mathsf{opt}(N_B)$ and $\mathsf{pess}(N_B)$: (a) The filled vertices are the points of S, and the hollow vertices are the points of S'. (b) The vertices of the set S_0. (c) The filled vertices are the points of S_1, and the hollow vertices are the points of S_2.

from the bottom-right (resp. bottom-left) vertices of the rectangles of N_B whose bottom sides are aligned with the bottom side of $\mathsf{MBR}(B)$; ℓ_1 (resp. ℓ'_1) a bottom-most vertex from the top-left (resp. bottom-left) vertices of the rectangles of N_B whose left sides are aligned with the left side of $\mathsf{MBR}(B)$; and ℓ_2 (resp. ℓ'_2) a top-most vertex from the bottom-left (resp. top-left) vertices of the rectangles of N_B whose left sides are aligned with the left side of $\mathsf{MBR}(B)$. All of these points can be found in $O(|N_B|)$ time, by a single pass over the elements of N_B (see Figure 8).

Suppose that $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a corner intersection. To compute $\mathsf{opt}(N_B)$, we need to compute the convex hull of the top-left vertices of the rectangles of N_B and the set V_B. Observe that such vertices that are not in the triangle T with vertex set $\{v_2, \ell'_2, t_2\}$ cannot be vertices of $\mathsf{opt}(N_B)$ (see Figure 8a). Then, we first find in $O(|N_B|)$ time the set S of the top-left vertices of the rectangles of N_B which belong to T, and after that compute $\mathsf{opt}(N_B) = conv(S \cup V_B)$ in $O(|S| \log |S|)$ time, using a standard algorithm for computing the convex hull. Doing this, we apply the convex hull algorithm for only the relevant set of points. Ideas similar to these ones are going to be used in the following. To compute $\mathsf{pess}(N_B)$, we find in $O(|N_B|)$ time the set S' of the bottom-left and top-right vertices of the rectangles of N_B that belong to the triangle with vertices $\{v_2, \ell_2, t_1\}$ (see Figure 8a), and then compute $\mathsf{pess}(N_B) = conv(S' \cup V_B)$ in $O(|S'| \log |S'|)$ time.

When $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a side intersection, we proceed as follows: To compute $\mathsf{opt}(N_B)$, we compute in $O(|N_B|)$ time the point set S_0 consisting of the top-left vertices of the rectangles of N_B that belong to the triangle with vertices $\{v_2, \ell_2', t_2\}$, and the bottom-left vertices that belong to the triangle with vertices $\{v_1, \ell_1', b_2\}$ (see Figure 8b). Then, we compute $\mathsf{opt}(N_B) = conv(S_0 \cup V_B)$ in $O(|S_0| \log |S_0|)$ time. To compute $\mathsf{pess}(N_B)$, we first find in $O(|N_B|)$ time the set S_1 containing the bottom-left and top-right vertices of the rectangles in N_B that belong to the triangle with vertices $\{v_2, \ell_2, t_1\}$, and the set S_2 containing the top-left and bottom-right vertices of the rectangles in N_B that belong to the triangle with vertices $\{v_1, \ell_1, b_1\}$ (see Figure 8c). After that, we compute the convex hulls $C_1 = conv(N_B' \cup V_B \cup \{v_1\}) = conv(S_1 \cup V_B \cup \{v_1\})$ and $C_2 = conv(N_B'' \cup V_B \cup \{v_2\}) = conv(S_2 \cup V_B \cup \{v_2\})$ in times $O(|S_1| \log |S_1|)$ and $O(|S_2| \log |S_2|)$, respectively. Finally, we compute $\mathsf{pess}(N_B)$ as the intersection $C_1 \cap C_2$. If $|C_1|$ and $|C_2|$ denote the numbers of vertices of C_1 and C_2, respectively, a representation of $C_1 \cap C_2$ can be computed in $O(\log |C_1| \cdot \log |C_2|)$ time in the worst case: If the point ℓ_1 is not above the point ℓ_2, then the boundary $C_1 \cap C_2$ consists of the clockwise boundary path L of C_2 connecting b_1 with ℓ_1, the segment connecting ℓ_1 with ℓ_2, the clockwise boundary path U of C_1 connecting ℓ_2 with t_1, and the similar path of C_1 that connects t_1 with b_1. Hence, in this case, a representation of $C_1 \cap C_2$ can be found in $O(1)$ time. Otherwise, if ℓ_1 is above ℓ_2, a representation of $C_1 \cap C_2$ is given by such above paths and the intersection point p between L and U. Observe that both L and U are x-monotone paths, and in each of them the edges are sorted from left to right. To find p, we need to find the edge of L that intersects U. This can be done with a binary search over the edges of L. Given any edge e of L, deciding whether e is to the left of U, intersects U, or is to the right of U, can be done by querying to which side of U is each endpoint of e (i.e. if the endpoint is or not inside C_1). Determining whether a given point belongs to a convex hull can be done with a binary search on the edges, running in $O(\log k)$ time, where k is the number of vertices. Then, each query costs $O(\log |C_1|)$ time, and $O(\log |C_2|)$ queries are performed, with a total running time of $O(\log |C_1| \cdot \log |C_2|)$.

4.2 Deciding convex hulls intersection

We show how to decide whether $= \mathsf{opt}(N_R)$ and $\mathsf{opt}(N_B)$ are disjoint. A similar procedure can be applied for $\mathsf{pess}(N_R)$ and $\mathsf{pess}(N_B)$.

Let $C_1 = \mathsf{opt}(N_R)$ and $C_2 = \mathsf{opt}(N_B)$. Suppose that $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a corner intersection. By Lemma 3, we can orient clockwise the edges of C_2 and consider the sequence S of consecutive edges that starts with the edge with source endpoint the bottom-left vertex of $\mathsf{MBR}(B)$, and ends with the edge with target endpoint the top-right vertex of $\mathsf{MBR}(B)$. The sequence S consists of three intervals: consecutive edges outside C_1 which point to C_1, consecutive edges which have at least one endpoint inside C_1, and consecutive edges outside C_1 which do not point to C_1. If the first edge of S does not point to C_1, then C_1 and C_2 are disjoint. In general, given any edge of C_2, querying to which interval of S the edge belongs to can be done with a binary search on the edges of C_1 from the top-right vertex to the bottom-left vertex clockwise, in $O(\log |C_1|)$ time. Furthermore, deciding whether there exists an edge of S that has at least one endpoint inside C_1 can be done with a binary search in S. The search performs $O(\log |C_2|)$ queries, and each query will cost $O(\log |C_1|)$ time. Deciding whether C_1 and C_2 are disjoint can thus be done in $O(\log |C_1| \cdot \log |C_2|)$ time. To find a line separating C_1 and C_2, in the case where they are disjoint, we need to find the first edge e in S that does not point to C_1. This edge can be found, similarly as above, in $O(\log |C_1| \cdot \log |C_2|)$ time. Let $\overline{e}$ denote the same edge e, but oriented in the contrary direction. If $\overline{e}$

does not point to C_1, then the straight line through e is a separating line. Otherwise, in $O(\log |C_1|)$ time we can find the edge e' of C_1 pointed by $\overline{e}$, and the line through e' is a separating line.

When $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a side intersection, a similar sequence S can be considered. In this case, S is the sequence of consecutive edges that starts with the edge with source endpoint the bottom-right vertex of $\mathsf{MBR}(B)$, and ends with the edge with target endpoint the top-right vertex of $\mathsf{MBR}(B)$. Deciding whether C_1 and C_2 are disjoint can be done in $O(\log |C_1| \cdot \log |C_2|)$ time.

4.3 Filtering rectangles

We show how to filter the rectangles of N_R and N_B, that is, to refine these sets by removing some elements, so that the new N_R and N_B still satisfy the properties (1) and (2) of Definition 4. We use the natural way of removing rectangles, say from N_B, which consists in removing all rectangles completely contained in the pessimistic convex hull (see the rectangle in dashed lines in Figure 6d and Figure 7). If the rectangle has a part outside the pessimistic convex hull, then it cannot be removed because such a part could contain blue points that are vertices of $conv(B \cup V_B)$.

Note that a rectangle N of N_B is contained in $\mathsf{pess}(N_B)$ if and only if the two left vertices of N belong to $\mathsf{pess}(N_B)$. Thus, once we have computed $\mathsf{pess}(N_B)$, determining whether N is contained in $\mathsf{pess}(N_B)$ can be done by querying twice whether a point belongs to $\mathsf{pess}(N_B)$. In this case, the two points are the two left vertices of N. Each query runs in $O(\log k)$ time, where k is the number of vertices of $\mathsf{pess}(N_B)$. The running time to filter the rectangles of N_B is then $O(|N_B| \cdot \log k) = O(|N_B| \cdot \log |N_B|)$. Symmetrically, a rectangle N of N_R is contained in $\mathsf{pess}(N_R)$ if and only if the two right vertices of N belong to $\mathsf{pess}(N_R)$, and similar decision and filtering algorithms can be used, where the filtering algorithm runs in $O(|N_R| \cdot \log |N_R|)$ time.

4.4 Separability algorithm

The algorithm consists of an outer procedure and an inner procedure. The outer procedure (see the pseudocode `DecideSeparability` of Figure 9) receives R and B as input, both represented in R-trees, and decide the linear separability of R and B. It also returns a separating line in the positive case. In this procedure, we first initialize the rectangle set N_R as the rectangles contained in the root node of the R-tree of R, and the rectangle set N_B as the rectangles contained in the root node of the R-tree of B. This allows to compute both $\mathsf{MBR}(R) = \mathsf{MBR}(N_R)$ and $\mathsf{MBR}(B) = \mathsf{MBR}(N_B)$. Then, we proceed as follows: If the intersection between $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ is empty, then we return a 'yes' answer together with an axis-aligned line containing an edge of $\mathsf{MBR}(R)$ that separates R and B. If $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a piercing intersection, then we return a 'no' answer. Otherwise, if $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a containment, corner, or side intersection, then the inner procedure (see the pseudocode `DecideSeparabilityCS` of Figure 10) is called accordingly. This procedure decides the linear separability when $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have corner or side intersection. Recall that when $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a containment intersection, the linear separation question can be reduced to solve (at most) four instances of the same question in which $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a corner intersection (see Section 2). This is done by extending the inner rectangle to contain one vertex of the outer rectangle. The inner procedure is as follows:

We start by computing the vertex sets V_R and V_B, according to the relative positions of $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$, which are necessary to the algorithm to distinguish between a corner and a side intersection. Then, the following actions with N_R and N_B are performed. We compute both $\mathsf{opt}(N_R)$

```
Algorithm DecideSeparability(R, B):
1.  N_R ← the set of rectangles in the root node of the R-tree of R
2.  N_B ← the set of rectangles in the root node of the R-tree of B
3.  MBR(R) ← MBR(N_R)
4.  MBR(B) ← MBR(N_B)
5.  if MBR(R) ∩ MBR(B) is empty then
6.      return true
7.  else if MBR(R) and MBR(B) have a piercing intersection then
8.      return false
9.  else if MBR(R) and MBR(B) have a containment intersection then
10.     if MBR(B) ⊂ MBR(R) then
11.         for each vertex v of MBR(R) do
12.             (* MBR(R) and MBR(MBR(B) ∪ {v}) have a corner intersection *)
13.             if DecideSeparabilityCS(N_R, N_B, MBR(R), MBR(MBR(B) ∪ {v})) then
14.                 return true
15.         return false
16.     else
17.         for each vertex v of MBR(B) do
18.             (* MBR(MBR(R) ∪ {v}) and MBR(B) have a corner intersection *)
19.             if DecideSeparabilityCS(N_R, N_B, MBR(MBR(R) ∪ {v}), MBR(B)) then
20.                 return true
21.         return false
22. else
23.     (* MBR(R) and MBR(B) have a corner or side intersection *)
24.     return DecideSeparabilityCS(N_R, N_B, MBR(R), MBR(B))
```

Figure 9: Algorithm to compute the linear separability of R and B.

and $\mathsf{opt}(N_B)$ (see Section 4.1), and test whether $\mathsf{opt}(N_R)$ and $\mathsf{opt}(N_B)$ are disjoint (see Section 4.2). If they are disjoint, then we report a 'yes' answer and find a separating line (see Section 4.2). Otherwise, if $\mathsf{opt}(N_R)$ and $\mathsf{opt}(N_B)$ are not disjoint, we continue as follows. We compute both $\mathsf{pess}(N_R)$ and $\mathsf{pess}(N_B)$, and decide whether they are disjoint (see sections 4.1 and 4.2). If the are not disjoint, then we report a 'no' answer. Otherwise, for each point set $X \in \{R, B\}$ such that N_X is made of rectangles, we filter N_X (see Section 4.3), and replace each remaining rectangle in N_X by its child rectangles, or points, in the corresponding R-tree. Observe that the new rectangles of N_X are all a level down to the level of the former rectangles in N_X. If at least one of the new N_R and N_B is made of rectangles (it can happen that one of N_R and N_B is made of rectangles and the other one is made of points since the R-trees of R and B can have different heights), then we repeat all this actions with these new N_R and N_B. Otherwise, if both N_R and N_B are made of points, we test whether $\mathsf{opt}(N_R) = \mathit{conv}(R)$ and $\mathsf{opt}(N_B) = \mathit{conv}(B)$ are disjoint, and find a separating line in the positive case, to finally decide whether R and B are linearly separable.

In the following, we analyse the asymptotic running time of the algorithm in the worst case. Let $m = |R|$ and $n = |B|$, and $r = O(\log m)$ and $b = O(\log n)$ denote the heights of the R-trees of R and B, respectively. Assume $r \leq b$ without loss of generality. For $i = 0, 1, \ldots, r$, let m_i denote the number of rectangles in the i-th level of the R-tree of R, and $N_R^{(i)}$ the set of rectangles N_R

Algorithm DecideSeparabilityCS($N_R, N_B, \mathsf{MBR}(R), \mathsf{MBR}(B)$):

```
1.  Compute V_R and V_B according to MBR(R) and MBR(B)
2.  opt_R ← OptimisticConvexHull(N_R, V_R)
3.  opt_B ← OptimisticConvexHull(N_B, V_B)
4.  while (true) do
5.      if opt_R and opt_B are disjoint then
6.          return true
7.      else if (N_R is made of points) and (N_B is made of points) then
8.          return false
9.      else
10.         if N_R is made of rectangles then
11.             pess_R ← PessimisticConvexHull(N_R, V_R)
12.         else
13.             pess_R ← opt_R
14.         if N_R is made of rectangles then
15.             pess_B ← PessimisticConvexHull(N_B, V_B)
16.         else
17.             pess_B ← opt_B
18.         if pess_R and pess_B are not disjoint then
19.             return false
20.         else
21.             if N_R is made of rectangles then
22.                 N_R ← Filter(N_R)
23.                 N_R ← ⋃_{N∈N_R} children(N)
24.                 opt_R ← OptimisticConvexHull(N_R, V_R)
25.             if N_B is made of rectangles then
26.                 N_B ← Filter(N_B)
27.                 N_B ← ⋃_{N∈N_B} children(N)
28.                 opt_B ← OptimisticConvexHull(N_B, V_B)
```

Figure 10: Algorithm to compute the linear separability of R and B when $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a corner or side intersection: $\mathsf{OptimisticConvexHull}(\cdot)$ returns the optimistic convex hull; $\mathsf{PessimisticConvexHull}(\cdot)$ returns the pessimistic convex hull; $\mathsf{Filter}(\cdot)$ removes from the argument the rectangles contained in the pessimistic convex hull; and $\mathsf{children}(N)$ returns the rectangles (or points) that are children of the rectangle N in the corresponding R-tree.

when N_R is formed by rectangles of the i-th level, where $|N_R^{(i)}| \leq m_i$ and $m_r = m$. Similarly, for $j = 0, 1, \ldots, b$, let n_j denote the number of rectangles in the j-th level of the R-tree of B, and $N_B^{(j)}$ the set of rectangles N_B when N_B is formed by rectangles of the j-th level, where $|N_B^{(j)}| \leq n_j$ and $n_b = n$. For level number $i = 0, 1, \ldots, r$, the algorithm in the worst case:

- computes both $\mathsf{opt}(N_R^{(i)})$ and $\mathsf{pess}(N_R^{(i)})$ in $O(|N_R^{(i)}| \cdot \log |N_R^{(i)}|)$ time;
- computes both $\mathsf{opt}(N_B^{(i)})$ and $\mathsf{pess}(N_B^{(i)})$ in $O(|N_B^{(i)}| \cdot \log |N_B^{(i)}|)$ time;
- decides $\mathsf{opt}(N_R^{(i)}) \cap \mathsf{opt}(N_B^{(i)}) = \emptyset$ in $O(\log |N_R^{(i)}| \cdot \log |N_B^{(i)}|)$ time;
- decides $\mathsf{pess}(N_R^{(i)}) \cap \mathsf{pess}(N_B^{(i)}) = \emptyset$ in $O(\log |N_R^{(i)}| \cdot \log |N_B^{(i)}|)$ time;
- filters $N_R^{(i)}$ in $O(|N_R^{(i)}| \cdot \log |N_R^{(i)}|)$ time, for $i < r$; and
- filters $N_B^{(i)}$ in $O(|N_B^{(i)}| \cdot \log |N_B^{(i)}|)$ time, for $i < b$.

For level number $j = r + 1, \ldots, b$, the algorithm in the worst case:

- computes both $\mathsf{opt}(N_B^{(j)})$ and $\mathsf{pess}(N_B^{(j)})$ in $O(|N_B^{(j)}| \cdot \log |N_B^{(j)}|)$ time;
- decides $\mathsf{opt}(N_R^{(r)}) \cap \mathsf{opt}(N_B^{(j)}) = \emptyset$ in $O(\log |N_R^{(r)}| \cdot \log |N_B^{(j)}|)$ time;
- decides $\mathsf{pess}(N_R^{(r)}) \cap \mathsf{pess}(N_B^{(j)}) = \emptyset$ in $O(\log |N_R^{(r)}| \cdot \log |N_B^{(j)}|)$ time; and
- filters $N_B^{(j)}$ in $O(|N_B^{(j)}| \cdot \log |N_B^{(j)}|)$ time, for $j < b$.

Summing up, the running time in the worst case is:

$$
\begin{aligned}
& \sum_{i=0}^{r} \left(O\left(|N_R^{(i)}| \cdot \log |N_R^{(i)}|\right) + O\left(|N_B^{(i)}| \cdot \log |N_B^{(i)}|\right) + O\left(\log |N_R^{(i)}| \cdot \log |N_B^{(i)}|\right)\right) + \\
& \sum_{i=r+1}^{b} \left(O\left(|N_B^{(i)}| \cdot \log |N_B^{(i)}|\right) + O\left(\log |N_R^{(r)}| \cdot \log |N_B^{(i)}|\right)\right) \\
= \; & \sum_{i=0}^{r} \Big(O(m_i \log m_i) + O(n_i \log n_i) + O(\log m_i \log n_i) \Big) + \sum_{i=r+1}^{b} \Big(O(n_i \log n_i) + O(\log m \log n_i) \Big) \\
= \; & \sum_{i=0}^{r} \Big(O(m_i \log m_i) + O(n_i \log n_i) \Big) + \sum_{i=r+1}^{b} \Big(O(n_i \log n_i) + O(\log m \log n_i) \Big) \\
= \; & O\left(\sum_{i=0}^{r} m_i \log m_i + \sum_{i=0}^{b} n_i \log n_i + \sum_{i=r+1}^{b} \log m \log n_i \right) \\
= \; & O\left(\log m \cdot \sum_{i=0}^{r} m_i + \log n \cdot \sum_{i=0}^{b} n_i + (b - r) \log m \log n \right)
\end{aligned}
$$

$$= \quad O\left(\log m \cdot \sum_{i=0}^{r} m_i + \log n \cdot \sum_{i=0}^{b} n_i + \log m \log^2 n\right).$$

Since an R-tree has the property that every node contains at least two children nodes, we have that $m_r = m$, $m_{r-1} \leq m/2$, $m_{r-2} \leq m/4$, $m_{r-3} \leq m/8$, and so on. That is, $m_i \leq m/2^{r-i}$ for $i = 0, 1, \ldots, r$. Similarly, $n_j \leq n/2^{b-j}$ for $j = 0, 1, \ldots, b$. The above running time is then:

$$\begin{aligned} O\left(\log m \cdot \sum_{i=0}^{r} m/2^i + \log n \cdot \sum_{i=0}^{b} n/2^i + \log m \log^2 n\right) &= O\left(m \log m + n \log n + \log m \log^2 n\right) \\ &= O(m \log m + n \log n). \end{aligned}$$

The worst case of our algorithm occurs when all nodes of the two R-trees, and all rectangles and points, need to be loaded to decide the linear separability of R and B. This happens in the following example. Suppose that all elements of B belong to the line $y = x$, for example, $B = \{(i, i) : i = 1, 2, \ldots, n\}$, and that $|R| = |B| = n$ with $R = \{(i - \varepsilon, i + \varepsilon) : i = 1, 2, \ldots, n\}$ for $\varepsilon = 1/2$. In this case, R and B are linearly separable, and $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$ have a corner intersection, where $\mathsf{MBR}(R)$ contains the top-left vertex of $\mathsf{MBR}(B)$, and $\mathsf{MBR}(B)$ contains the bottom-right vertex of $\mathsf{MBR}(R)$. Consider any step of our algorithm (refer to Figure 10), with N_R and N_B representing the rectangles from the R-trees of R and B, respectively. Observe that in every rectangle of N_B the diagonal connecting the bottom-left vertex with the top-right vertex is contained in the line $y = x$. This implies that the pessimistic convex hull $\mathsf{pess}(N_B)$ equals the triangle with vertices $(1, 1)$, $(n, 1)$, and (n, n), and that no rectangle of N_B is contained in $\mathsf{pess}(N_B)$. A similar situation occurs with N_R: the pessimistic convex hull $\mathsf{pess}(N_R)$ equals the triangle with vertices $(1 - \varepsilon, 1 + \varepsilon)$, $(1 - \varepsilon, n + \varepsilon)$, and $(n - \varepsilon, n + \varepsilon)$, and no rectangle of N_R is contained in $\mathsf{pess}(N_R)$. Then, no rectangle of N_R or N_B can be discarded in any step of the algorithm. Furthermore, $\mathsf{pess}(N_R)$ and $\mathsf{pess}(N_B)$ are disjoint, whereas the optimistic convex hulls $\mathsf{opt}(N_R)$ and $\mathsf{opt}(N_B)$ are not disjoint (due to the way we choose ε). All of these observations imply that the algorithm will stop alfter loading all points, hence all nodes and rectangles, from both R-trees.

4.5 The convex hull of a point set

Let P be a finite point set in the plane, given in an R-tree. In this section, we present an algorithm to compute $conv(P)$. Let N_P be a set of rectangles from the R-tree of P such that N_P satisfies $conv(P) \subseteq conv(N_P)$ (similar as Definition 4). To compute $conv(P)$, we first discard rectangles N from N_P such that all points of P contained in N are not vertices of $conv(P)$. After that, we replace each rectangle N that remains in N_P by its child rectangles (or points) in the R-tree. We stop when all elements of N_P are points, and return $conv(N_P)$. Let v_1, v_2, v_3, and v_4 be the top-left, bottom-left, bottom-right, and top-right vertices of $\mathsf{MBR}(P)$, respectively. To discard rectangles from N_P, we compute the following convex hulls by following ideas similar to those given in Section 4.1: $C_1 = conv(\{\mathtt{SE}(N) \mid N \in N_P\} \cup \{v_2, v_3, v_4\})$, $C_2 = conv(\{\mathtt{NE}(N) \mid N \in N_P\} \cup \{v_1, v_3, v_4\})$, $C_3 = conv(\{\mathtt{NW}(N) \mid N \in N_P\} \cup \{v_1, v_2, v_4\})$, and $C_4 = conv(\{\mathtt{SW}(N) \mid N \in N_P\} \cup \{v_1, v_2, v_3\})$. Given a rectangle N of N_P, if N is contained in the intersection $C_1 \cap C_2 \cap C_3 \cap C_4$, then no point of P contained in N can be a vertex of $conv(P)$. Furthermore, N is contained in $C_1 \cap C_2 \cap C_3 \cap C_4$ if and only if the top-left, bottom-left, bottom-right, and top-right vertices of N are contained in C_1, C_2, C_3, and C_4, respectively. Once such four convex hulls are computed, these four decisions can

Algorithm `ConvexHull`(P):
1. $N_P \leftarrow$ the set of rectangles in the root node of the R-tree of P
2. $\mathsf{MBR}(P) \leftarrow \mathsf{MBR}(N_P)$
3. **repeat**
4. $\quad C_1 \leftarrow conv(\{\mathtt{SE}(N) \mid N \in N_P\} \cup \{v_2, v_3, v_4\})$
5. $\quad C_2 \leftarrow conv(\{\mathtt{NE}(N) \mid N \in N_P\} \cup \{v_1, v_3, v_4\})$
6. $\quad C_3 \leftarrow conv(\{\mathtt{NW}(N) \mid N \in N_P\} \cup \{v_1, v_2, v_4\})$
7. $\quad C_4 \leftarrow conv(\{\mathtt{SW}(N) \mid N \in N_P\} \cup \{v_1, v_2, v_3\})$
8. $\quad N_P \leftarrow N_P \setminus \{N \in N_P \mid N \subset C_1 \cap C_2 \cap C_3 \cap C_4\}$
9. $\quad N_P \leftarrow \bigcup_{N \in N_P} \mathsf{children}(N)$
10. **until** N_P is made of points
11. **return** $conv(N_P)$

Figure 11: Algorithm to compute the convex hull of a point set P given in an R-tree.

be made in times $O(\log |C_1|)$, $O(\log |C_2|)$, $O(\log |C_3|)$, and $O(\log |C_4|)$, respectively. The algorithm to compute $conv(P)$ is described in the pseudocode of Figure 11. The running time is $O(n \log n)$, where n is the number of points, and can be obtained from arguments similar to those of Section 4.4.

5 Experimental results

In this section, we describe the experiments that we implemented to evaluate the performance of our separability testing algorithm in terms of running time (via counting the number of access to nodes of the R-trees) and memory usage. The algorithm was implemented in the `C++` language, using the implementation of the R-tree data structure of the library `LibSpatialIndex` [10]. Nodes of size 1K were used to build the R-trees. The experiments were executed in a `Lenovo ThinkPad x240` computer, with 8GB of `RAM` memory, and an `Intel Core i5 4300U` microprocessor, and both real and synthetic data were considered.

5.1 Real data

We consider a first data set consisting of 200 thousands of MBRs representing spatial objects of California roads, and a second data set consisting of 2.2 millions of MBRs representing spatial objects of rivers of Iowa, Kansas, Missouri, and Nebraska [3]. From the first set we generate a set of red points by taking from each MBR its center point, and perform a similar operation to the second data set to obtain a set of blue points. To test our algorithm, we mapped both point sets to the space $[0, 1] \times [0, 1]$. Each colored point set was stored in a different R-tree, and the numbers of blocks used in the R-trees are shown in Table 1. In Figure 12a, we make a graphic representation of both colored point sets, where we plot only about the 10% of the points of each set. In Figure 12b, we draw the MBR of each colored point set, showing that they have a side intersection with a considerably high overlapping, precisely, the common area of the MBRs is above the 98% of the total area. In Table 1 we show the results of the execution of our linear separability testing algorithm: a 2.79% of the nodes (i.e. memory blocks) of the red R-tree are

accessed, whereas a 1.36% of the nodes of the blue R-tree are accessed. To solve this particular instance of the problem, only 40Kb of main memory is required.

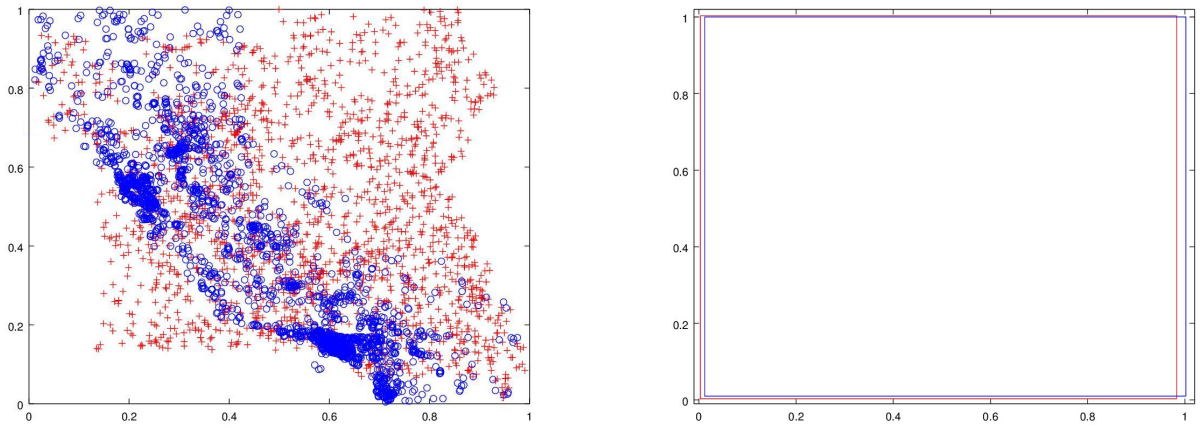

(a) (b)

Figure 12: a) Colored point sets generated from real data. Each red points is represented by a +, and each blue point by a ∘. b) The MBR of each point set, where the MBRs has a side intersection and their intersection area is above the 98% of the total area.

	size	% nodes accessed	size of R-tree (# of blocks)
Red points	200,000	2.79%	12,178
Blue points	2,200,000	1.36%	35,965

Table 1: Results of experiments on real data.

5.2 Synthetic data

We extecute our algorithm on several synthetic data sets, each data set consisting of colored point sets randomly generated in the range $[0,1] \times [0,1] = [0,1]^2$. To generate a point set $R \cup B$, we proceed in the following steps:

1. We pick the number of points that each color class will contain. This number is either 1, 2, 5, or 10 millions.

2. We define two rectangles $R_r, R_b \subset [0,1]^2$ of equal areas, so that R and B will be generated inside R_r and R_b, respectively, and the area of $R_r \cap R_b$ represents a given percent of the areas of R_r and R_b. We select such a percent among 1%, 5%, 10%, and 50%. Furthermore, we also fix the type of intersection of R_r and R_b (hence the type of intersection of $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$): corner or side.

3. We select the distribution in which each point set (R and B) is generated inside its corresponding rectangle. We consider two distributions: uniform and Gaussian.

Size of point sets (in millions)	Red R-tree (disk blocks)	Blue R-tree (disk blocks)	Total (disk blocks)
1	61,386	61,388	122,774
2	122,210	122,226	244,436
5	304,494	304,858	609,352
10	608,671	609,749	1,218,420

Table 2: Number of blocks (i.e. nodes) used in average to build the R-trees.

Size of point sets (in millions)	Type of intersection							
	Corner				Side			
	% of intersection				% of intersection			
	1%	5%	10%	50%	1%	5%	10%	50%
1	0.54	0.31	0.41	0.34	0.34	1.14	1.08	2.11
2	0.23	0.13	0.09	0.35	0.22	0.39	0.58	0.19
5	0.08	0.11	0.03	0.05	0.12	0.36	0.28	0.36
10	0.18	0.03	0.03	0.44	0.09	0.18	0.2	0.27

Table 3: Percentage of the nodes of the R-trees that were accessed, in the cases where the points were generated by using the uniform distribution.

In total, we run our algorithm on 64 synthetic data sets, accounting from 4 possibilities for the generated number of points, times 4 percents of intersection area, times 2 types of intersection of $\mathsf{MBR}(R)$ and $\mathsf{MBR}(B)$, times 2 distributions. In Figure 13, we show examples of colored point sets generated with a Gaussian distribution having their MBRs a corner intersection.

For each possible number of points in the generated point sets, the average size of the R-trees, expressed as the number of disk blocks (i.e. nodes), is shown in Table 2. We also measured the percentage of the nodes of the R-trees that were accessed by our algorithm, as shown in Figure 14, Table 3, and Table 4.

Size of point sets (in millions)	Type of intersection							
	Corner				Side			
	% of intersection				% of intersection			
	1%	5%	10%	50%	1%	5%	10%	50%
1	0.45	0.13	0.41	0.13	0.41	0.16	0.29	0.01
2	0.25	0.24	0.23	0.17	0.18	0.08	0.09	0.19
5	0.03	0.03	0.03	0.03	0.08	0.00	0.01	0.00
10	0.06	0.05	0.06	0.05	0.01	0.00	0.01	0.00

Table 4: Percentage of the nodes of the R-trees that were accessed, in the cases where the points were generated by using the Gaussian distribution.

The results indicate that when points are generated by a uniform distribution, the number of R-tree nodes accessed increases as the percentage of intersection between R_r and R_b is increased (see Figure 14a and Figure 14b). For example, for 1 million of points generated and a 50% of area of

Size of point sets (in millions)	Type of intersection							
	Corner				Side			
	% of intersection				% of intersection			
	1%	5%	10%	50%	1%	5%	10%	50%
1	24	24	25	26	26	27	27	29
2	42	41	46	46	44	43	43	43
5	9	11	9	11	14	12	14	12
10	17	16	17	21	18	19	18	17

Table 5: Memory usage (in Kb) in sets with uniform distribution.

intersection, and a side intersection, it is needed to access to a 2.11% of the nodes of the R-trees, whereas under the same conditions but a 1% of intersection, it is needed to access a 0.34% of the nodes. For point sets generated with the Gaussian distribution (Figure 14c and Figure 14d), the percentage of nodes accessed seems to not depend on the percentage of intersection. For example, for 1 million of points generated, if the % of area of intersection between the MBRs is 1.0% or 10.0%, it is needed to access a 0.45% or 0.41% of the nodes, respectively. In Figures 14a, 14b, 14c, and 14d, we can note that when the number of points generated is increased, the percentage of nodes accessed decreases, being in all cases less that 0.1%.

In Table 5 and Table 6, we note the amount of memory used by the algorithm for the point sets with uniform and Gaussian distributions. This includes the amount of memory used for the node lists and convex hulls (optimistic and pessimistic). We can also note that the required ranges between 15Kb and 47Kb. Furthermore, the memory is not affected when we increase the side of the point sets generated. In Figure 15, it is shown that the amount of memory required has a similar behavior, independently of the type of intersection of the MBRs and the type of distribution.

Size of point sets (in millions)	Type of intersection							
	Corner				Side			
	% of intersection				% of intersection			
	1%	5%	10%	50%	1%	5%	10%	50%
1	26	23	25	23	33	26	28	25
2	44	44	44	43	47	45	43	43
5	8	8	9	8	15	7	9	7
10	19	17	17	17	19	20	17	14

Table 6: Memory usage in sets with Gaussian distribution.

6 Conclusions

In this paper, we have proposed an algorithm to decide the linear separability of two point sets of cardinalities n and m, respectively, both sets stored in a different R-tree. The algorithm takes advantage of the properties of the R-trees in order to access as less nodes as possible. The running time complexity in the worst case is within $O(m \log m + n \log n)$. With the goal of evaluating the performance of the algorithm in practice, we designed several experiments with both real and

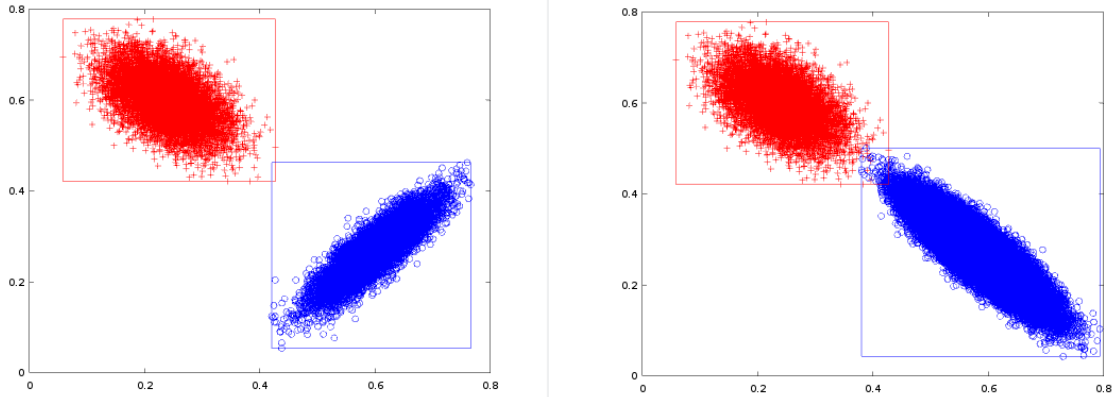

(a) Corner, 1% of intersection.

(b) Corner, 5% of intersection.

Figure 13: Synthetic data with a Gaussian distribution.

synthetic point sets, and an implementation of the algorithm was run in each experiment. The results of the experiments showed that the algorithm performs few accesses to disk (i.e. accesses to nodes of the R-trees), uses a small amount of RAM memory and a low computation time.

Our algorithm expands the use of the R-trees, a multidimensional data structure well used in several spatial database systems such as Postgres and Oracle. According to the bibliography review, and to the best of our knowledge, this is the first algorithm that tackles the geometric separability of massive spatial object sets stored in secondary storage data structures.

For future work, we propose the study of other types of geometric separability problems when the input is given in R-trees or other secondary-storage spatial data structures, for example separating red and blue points by axis parallel rectangles, wedges, or constrained polylines. We also propose to design an extension of this algorithm to work in dimensions higher that two.

Acknowledgements

C. T. was supported by CONICYT scholarship PCHA/MagisterNacional/2015-22151665 of the Government of Chile and the research group Bases de Datos 132019 GI/EF funded by Universidad del Bío-Bío (Chile).
P. P-L. was supported by project Millennium Nucleus Information and Coordination in Networks ICM/FIC RC130003 (Chile).
G. G. was supported by the research group Bases de Datos 132019 GI/EF, and the research project Diseño e Implementación de Algoritmos Geométricos en el contexto de Bases de Datos Espaciales 142719 3/R, both funded by Universidad del Bío-Bío (Chile).

References

[1] K. P. Bennett and E. J. Bredensteiner. Duality and geometry in SVM classifiers. In *Proceedings of the Seventeenth International Conference on Machine Learning*, ICML '00, pages 57–64, San Francisco,

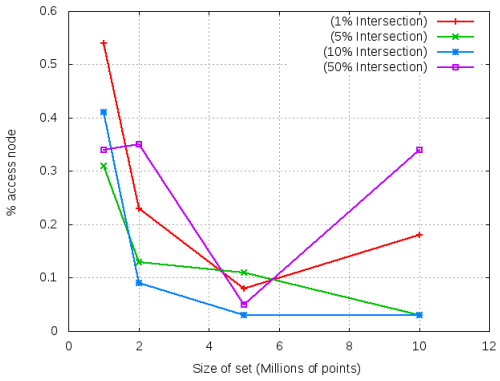

(a) Sets with uniform distribution and corner intersection.

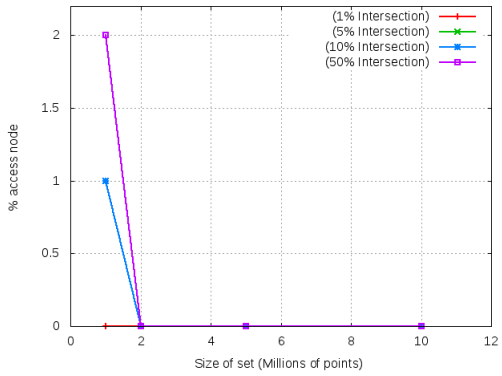

(b) Sets with uniform distribution and side intersection.

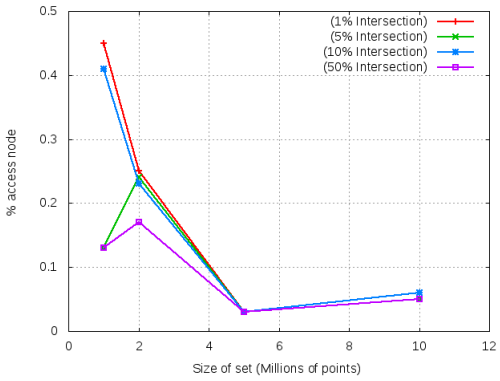

(c) Sets with Gaussian distribution and corner intersection.

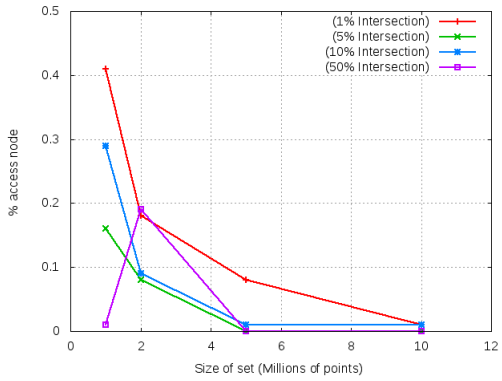

(d) Sets with Gaussian distribution and side intersection.

Figure 14: Percentage of nodes accessed.

CA, USA, 2000. Morgan Kaufmann Publishers Inc.

[2] C. Böhm and H.-P. Kriegel. Determining the convex hull in large multidimensional databases. In *Proceedings of the Third International Conference on Data Warehousing and Knowledge Discovery*, DaWaK '01, pages 294–306, London, UK, UK, 2001. Springer-Verlag.

[3] Chorochronos. Datasets & Algorithms. `http://www.rtreeportal.org`. Accessed: 2015-11-30.

[4] A. Corral, Y. Manolopoulos, Y. Theodoridis, and M. Vassilakopoulos. Algorithms for processing k-closest-pair queries in spatial databases. *Data Knowl. Eng.*, 49(1):67–104, Apr. 2004.

[5] J. Czyzowicz, E. Rivera-Campo, J. Urrutia, and J. Zaks. Separating convex sets in the plane. *Discrete Comput. Geom.*, 7(2):189–195, 1992.

[6] M. E. Dyer. Linear time algorithms for two- and three-variable linear programs. *SIAM J. Comput.*, 13(1):31–45, 1984.

[7] V. Gaede and O. Günther. Multidimensional access methods. *ACM Comput. Surv.*, 30(2):170–231, June 1998.

[8] G. Gutiérrez, J. R. Paramá, N. R. Brisaboa, and A. Corral. The largest empty rectangle containing only a query object in spatial databases. *GeoInformatica*, 18(2):193–228, 2014.

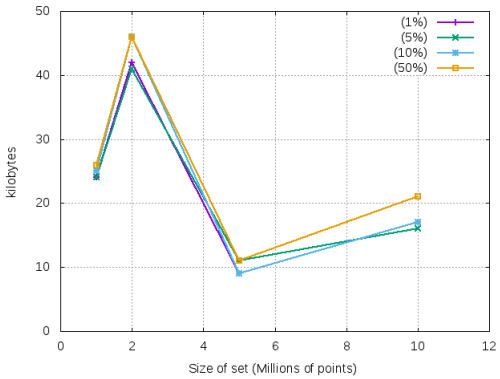

(a) Uniform distribution, corner intersection.

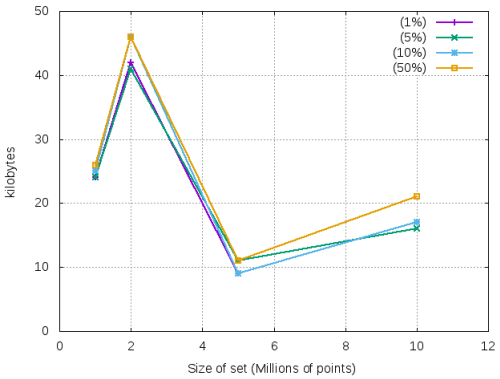

(b) Uniform distribution, side intersection.

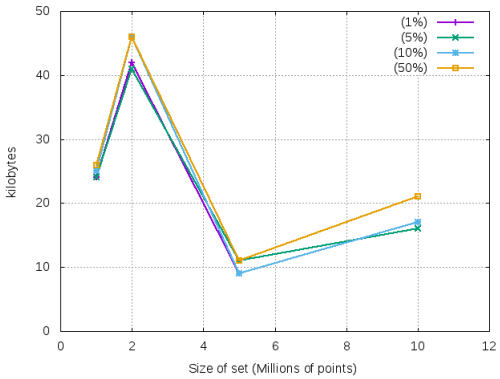

(c) Gaussian distribution, corner intersection.

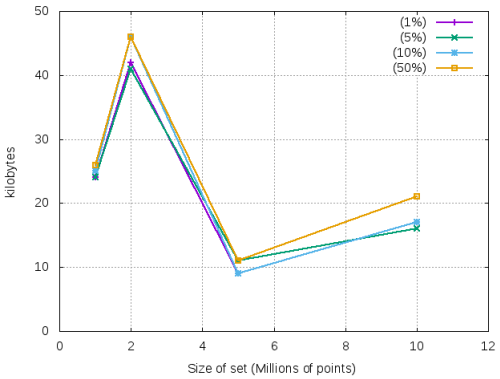

(d) Gaussian distribution, side intersection.

Figure 15: Memory usage.

[9] A. Guttman. R-trees: A dynamic index structure for spatial searching. *SIGMOD Rec.*, 14(2):47–57, June 1984.

[10] M. Hadjieleftheriou. LibSpatialIndex. `http://libspatialindex.org/`. Accessed: 2015-11-30.

[11] D. J. Hand, P. Smyth, and H. Mannila. *Principles of Data Mining*. MIT Press, Cambridge, MA, USA, 2001.

[12] G. R. Hjaltason and H. Samet. Incremental distance join algorithms for spatial databases. *SIGMOD Rec.*, 27(2):237–248, June 1998.

[13] Y. Manolopoulos, A. Nanopoulos, A. N. Papadopoulos, and Y. Theodoridis. *R-Trees: Theory and Applications (Advanced Information and Knowledge Processing)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.

[14] N. Megiddo. Linear programming in linear time when the dimension is fixed. *J. ACM*, 31(1):114–127, 1984.

[15] N. Roussopoulos, S. Kelley, and F. Vincent. Nearest neighbor queries. *SIGMOD Rec.*, 24(2):71–79, May 1995.

[16] R. Seidel. Small-dimensional linear programming and convex hulls made easy. *Discrete Comput. Geom.*, 6:423–434, 1991.

[17] S. Shekhar and S. Chawla. *Spatial databases - a tour.* Prentice Hall, 2003.

[18] G. T. Toussaint. A simple linear algorithm for intersecting convex polygons. *The Visual Computer*, 1(2):118–123, 1985.